arXiv:1612.05179v2 [stat.ME] 16 Dec 2016

Regression assisted inference for the average treatment effect in paired experiments

Colin B. Fogarty *

Abstract

In paired randomized experiments individuals in a given matched pair may differ on prognostically important covariates despite the best efforts of practitioners. We examine the use of regression adjustment as a way to correct for persistent covariate imbalances after randomization, and present two regression assisted estimators for the sample average treatment effect in paired experiments. Using the potential outcomes framework, we prove that these estimators are consistent for the sample average treatment effect under mild regularity conditions even if the regression model is improperly specified. Further, we describe how asymptotically conservative confidence intervals can be constructed. We demonstrate that the variances of the regression assisted estimators are at least as small as that of the standard difference-in-means estimator asymptotically. Through a simulation study, we illustrate the appropriateness of the proposed methods in small and moderate samples. The analysis does not require a superpopulation model, a constant treatment effect, or the truth of the regression model, and hence provides a mode of inference for the sample average treatment effect with the potential to yield substantial improvements in the power of the resulting analysis over the classical analysis without imposing potentially unrealistic assumptions.

1 Introduction

The use of paired experiments has historically been limited to cases in which pairs are formed using a small number of binary or categorical covariates. In practice, there are often continuous covariates which are also believed to be predictive for the potential outcomes under treatment and control. Greevy et al. (2004) introduced a form of multivariate matching before randomization which assigns a group of $2n$ experimental units to n pairs to minimize the within-pairs covariate distance, hence improving covariate balance for many variables at the same time. The pairs are arrived upon through the solution of a non-bipartite matching problem. Once the pairs are established, exactly one subject in each pair is randomly assigned to the treatment. Fisherian inference for no treatment effect and Neymanian inference for the average treatment effect can then proceed with respect to the randomization distribution generated by this paired design; see, for example, Rosenbaum (2002) and Imai (2008). Greevy et al. (2004) demonstrate the improvements produced by this procedure in terms of both covariate balance and power for the resulting inference.

*Operations Research and Statistics Group, MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge MA 02142 (e-mail: cfogarty@mit.edu)

While mitigable through the method of Greevy et al. (2004), persistent within-pair discrepancies on the basis of continuous covariates of interest are unavoidable. These discrepancies may, in turn, yield chance imbalances on the basis of these covariates in any given randomization. One common strategy for accounting for remaining imbalances is through covariance adjustment. Freedman (2008) and Lin (2013) investigates the impact of regression adjustment on Neymanian inference for the average treatment effect in completely randomized experiments. In the experimental setup considered by Neyman, the estimand of interest is the average treatment effect for the individuals in the sample at hand (Neyman, 1923). Neither correctness of the fitted regression model nor a superpopulation model (that units are drawn independently and identically distributed from a larger population) are assumed. Instead, inference proceeds using the physical act of randomization as its sole justification, and asymptotic calculations consider sequences of experiments of increasing size without specifying how the experimental units were sampled. In this context, Lin (2013) discusses how regression adjustment can yield a consistent estimator of the average treatment effect whose asymptotic variance is no larger than the standard difference of means estimator regardless of the truth of the underlying model. The performed inference is then agnostic to the truth of the underlying model and instead leverages the adjustment as an algorithmic fit in an attempt to yield a more efficient estimator. See also Aronow and Middleton (2013) and Bloniarz et al. (2016) for related work with completely randomized experiments.

In the context of paired randomized experiments, Rosenbaum (2002) describes how covariance adjustment can be utilized to yield exact tests under the assumption of a constant treatment effect through the use of randomization inference. Imbens and Rubin (2015, §10) discuss how under a superpopulation model, regression can be utilized to yield a consistent estimator for the average treatment effect in the superpopulation. In what follows, we show that covariance adjustment can also be leveraged for inference on the sample average treatment effect in a paired experiment in a manner that is agnostic to the truth of the fitted model.

2 A paired randomized experiment with covariates

There are n independent matched pairs. In each of the i pairs there is one individual who is randomly assigned the treatment, denoted as $Z_{ij} = 1$, and one who receives the control, denoted as $Z_{ij'} = 0$, such that $Z_{i1} + Z_{i2} = 1$. Individual j in matched set i has a P-dimensional vector of measured covariates $x_{ij} = (x_{ij1}, ..., x_{ijP})$. Each individual has a potential outcome under treatment, r_{Tij}, and under control, r_{Cij}, $i = 1, ..., n; j = 1, 2$. The pair of potential outcomes (r_{Tij}, r_{Cij}) is not jointly observable for any individual. Instead, we observe the response $R_{ij} = r_{Tij}Z_{ij} + r_{Cij}(1 - Z_{ij})$ for each individual. As a consequence, the individual level treatment effect $\tau_{ij} = r_{Tij} - r_{Cij}$ is not observable for any individual (Neyman, 1923; Rubin, 1974).

Let Ω be the set of 2^n possible values of $Z = (Z_{11}, Z_{12}, ..., Z_{n2})$ under the matched pairs design, namely those where $Z_{i1} + Z_{i2} = 1$ for $i = 1, ..., n$. Each $z \in \Omega$ has probability 2^{-n} of being selected; let $\mathcal{Z}$ denote the event $Z \in \Omega$. Quantities dependent on the assignment vector such as Z and $R = (R_{11}, R_{12}, ..., R_{n2})$ are random, whereas $\mathcal{F} = \{(r_{Tij}, r_{Cij}, x_{ij}), i = 1, ..., n, j = 1, 2\}$ contains fixed quantities. In a paired experiment, $\text{pr}(Z = z \mid \mathcal{F}, \mathcal{Z}) = \text{pr}(Z = z \mid \mathcal{Z}) = 2^{-n}$, and $\text{pr}(Z_{ij} = 1 \mid \mathcal{F}, \mathcal{Z}) = \text{pr}(Z_{ij} = 1 \mid \mathcal{Z}) = 1/2$ for $i = 1, ..., n$, $j = 1, 2$.

Write $d_i = f(x_{i1}) - f(x_{i2})$ for some function $f : \mathbb{R}^P \to \mathbb{R}^K$ with K fixed as the difference

in transformed covariates between unit 1 and unit 2 in matched pair i. Similarly write $m_i = \{g(x_{i1})+g(x_{i2})\}/2-n^{-1}\sum_{\ell=1}^{n}\{g(x_{\ell 1})+g(x_{\ell 2})\}/2$ for some function $g:\mathbb{R}^P \to \mathbb{R}^L$, as the average of the transformed covariates in matched pair i relative to the mean across matched pairs. Write D for the $n\times K$ matrix whose ith row consists of d_i^T, and write M for the $n\times L$ matrix whose ith row consists of m_i^T. In what follows we treat K and L as fixed and require that $K+L+1<n$.

Let $H_D = D(D^TD)^{-1}D$ be the "hat matrix" for D, i.e. the orthogonal projection of $\mathbb{R}^K$ onto the column space of D, and let H_M be that hat matrix for M. Define $V_i = 2Z_{i1}-1$, such that $E(V_i \mid \mathcal{F},\mathcal{Z})=0$, and let V be the $n\times n$ matrix with V_i on the ith diagonal and zeroes on the off-diagonal. The matrix consisting of the differences between the transformed covariates of the treated individuals minus those of the control individuals can then be written as VD, and the associated projection matrix onto the columnspace of VD is $VD(D^TVVD)^{-1}DV = VH_DV$ since $V_i^2=1$. Finally, write A for the $n\times(K+L)$ matrix whose ith row is $a_i^T = (V_id_{i1},...,V_id_{iK},m_{i1},...,m_{iL})^T$, and write H_A for the hat matrix corresponding to A.

3 The sample average treatment effect

Let $\varphi_{i1} = r_{Ti1}-r_{Ci2}$ and $\varphi_{i2} = r_{Ti2}-r_{Ci1}$ be the observed paired difference if $Z_{i1}=1$ and $Z_{i2}=1$ respectively. The treated minus control difference in pair i that is actually observed is

$$Y_i = Z_{i1}\varphi_{i1} + Z_{i2}\varphi_{i2}.$$

Write $\varphi_j = (\varphi_{1j},...,\varphi_{nj})$ for $j=1,2$, and write $Y=(Y_1,...,Y_n)$. The average treatment effect in a paired experiment with n pairs is

$$\bar{\tau}_n = \frac{1}{2n}\sum_{i=1}^{n}(\tau_{i1}+\tau_{i2}) = \frac{1}{2n}\sum_{i=1}^{n}(\varphi_{i1}+\varphi_{i2}).$$

Henceforth we will write $\bar{\tau}=\bar{\tau}_n$, but the dependence of this and other causal estimands on n should be kept in mind. The classical estimator for $\bar{\tau}$ in a paired experiment is the observed average of paired differences $\hat{\tau}_C = \bar{Y} = n^{-1}\sum_{i=1}^{n}Y_i$. Imai (2008) demonstrates that $E(\hat{\tau}_C \mid \mathcal{F},\mathcal{Z}) = \bar{\tau}$. The variance for $\hat{\tau}$ across randomizations is $\text{var}(\hat{\tau}_C \mid \mathcal{F},\mathcal{Z}) = (1/4)n^{-2}\sum_{i=1}^{n}(\varphi_{i1}-\varphi_{i2})^2$, which is unknown in practice because it depends on the missing potential outcomes. Imai (2008) shows that the classical estimator of the variance of the difference in means in a paired study, $S_C^2 = \{n(n-1)\}^{-1}\sum_{i=1}^{n}(Y_i-\hat{\tau}_C)^2$, is always an upper bound on $\text{var}(\hat{\tau}_C \mid \mathcal{F},\mathcal{Z})$ in expectation. Furthermore, under mild regularity conditions $nS_C^2 - \text{var}(n^{1/2}\hat{\tau}_C \mid \mathcal{F},\mathcal{Z})$ converges to a nonnegative value in probability.

To test the null $H_0:\bar{\tau}=\bar{\tau}_0$, inference proceeds under a normal approximation and using S_C^2 in place of $\text{var}(\hat{\tau}_C \mid \mathcal{F},\mathcal{Z})$. That is, under the null $\bar{\tau}=\bar{\tau}_0$ we approximate $\text{pr}(\hat{\tau}_C \le t)$ by $\Phi\{(t-\bar{\tau}_0)/S_C\}$, where $\Phi(\cdot)$ is the standard normal cumulative distribution function. This normal approximation is justified under regularity conditions on the sequence of constants $\{(\varphi_{i1},\varphi_{i2}), i=1,...,n\}$ as $n\to\infty$; see Imai (2008) and Ding (2016) among many for more details.

4 Regression assisted estimators

4.1 Two intercepts from multivariate regressions

We consider the intercept coefficients from a regression of Y on VD, $\hat{\tau}_{R1}$, and of Y on VD and M, $\hat{\tau}_{R2}$, as estimators for $\bar{\tau}$. It will be useful to express these quantities as

$$\hat{\tau}_{R1} = \frac{e^T(I - VH_DV)Y}{e^T(I - VH_DV)e}; \qquad \hat{\tau}_{R2} = \frac{e^T(I - H_A)Y}{e^T(I - H_A)e},$$

where I is the identity matrix, and e is the vector containing n ones. The numerators are simply the sum of the residuals from regressions without a constant column, while the denominator is the sum of residuals from regressions with e as the dependent variable (again without a constant column). As discussed in Imbens and Rubin (2015), a regression on VD encodes a belief that the difference between covariates in matched pair i, d_i, may be predictive of $\varphi_{i1} - \varphi_{i2}$. A regression including both VD and M encodes a belief that the difference in covariates in matched pair i is predictive of $\varphi_{i1} - \varphi_{i2}$, and that the relative level of the covariates m_i is predictive of $\varphi_{i1} + \varphi_{i2} = \tau_{i1} + \tau_{i2}$.

Under an additive treatment effect model, i.e. $r_{Tij} = r_{Cij} + \tau$ for $i = 1, ..., n$; $j = 1, 2$, the relative level of the covariates m_i cannot be predictive of the aggregate treatment effect $\tau_{i1} + \tau_{i2}$ in matched pair i as the aggregate treatment effect is constant across matched pairs by assumption. Hence, there is nothing to be gained by including M in the regression. In such a model, $\hat{\tau}_{R1}$ can be motivated as a Hodges-Lehmann estimator based on the test-statistic $T(\tau_0) = e^T(I - VH_DV)(Y - \tau_0 e)$ for the null hypothesis $\tau = \tau_0$ (Hodges and Lehmann, 1963; Rosenbaum, 1993), as $E(T(\tau_0) \mid \mathcal{F}, \mathcal{Z}) = 0$ under $\tau = \tau_0$ and $\hat{\tau}_{R1}$ is the value for a that solves $T(a) = 0$. Under this model, $\hat{\tau}_{R1}$ is an unbiased estimator for τ, and inference for τ can be conducted using the randomization distribution of $T(\tau_0)$ as discussed in Rosenbaum (2002).

4.2 Limiting distributions

Without the assumption of additivity, $\hat{\tau}_{R1}$ is no longer unbiased for $\bar{\tau}$. Further, the relative level of the covariates M may be predictive of the treatment effects in a given matched pair, potentially motivating the use of the estimator $\hat{\tau}_{R2}$. We now demonstrate that both $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ can be used to facilitate inference on $\bar{\tau}$ without assuming an additive treatment effect model. Moving forward, we impose two regularity conditions.

Condition 1. *(Bounded Fourth Moments)* . There exists a $C < \infty$ such that, for all n, $n^{-1}\sum_{i=1}^n(\varphi_{i1} + \varphi_{i2})^4 < C$, $n^{-1}\sum_{i=1}^n(\varphi_{i1} - \varphi_{i2})^4 < C$, $n^{-1}\sum_{i=1}^n d_{ik}^4 < C$ for $k = 1, .., K$, and $n^{-1}\sum_{i=1}^n m_{i\ell}^4 < C$ for $\ell = 1, .., C$.

Condition 2. *(Cesàro Summability).* $n^{-1}\sum_{i=1}^n(\varphi_{i1}+\varphi_{i2})$, $n^{-1}\sum_{i=1}^n(\varphi_{i1}+\varphi_{i2})^2$, and $n^{-1}\sum_{i=1}^n(\varphi_{i1}-\varphi_{i2})^2$ converge to finite limits as $n \to \infty$. Further, $n^{-1}\sum_{i=1}^n(\varphi_{i1} - \varphi_{i2})d_{ik}$ and $n^{-1}\sum_{i=1}^n(\varphi_{i1} + \varphi_{i2})m_{i\ell}$ converge to a finite limit for $k = 1, ..., K$ and $\ell = 1, ..., L$ as $n \to \infty$. Finally, $n^{-1}A^TA$ converges to a finite, invertible matrix.

Remark 1. Condition 1 imposes boundedness of fourth moments and mirrors Condition 9 (b) of Freedman (2008) and Condition 1 of Lin (2013). Condition 2 ensures the existence of a limit for

the regression coefficients and limits of various population means, resembling Condition 2 of Lin (2013).

Theorem 1. *Define* $\hat{\tau}_{R^*}$ *as*

$$\hat{\tau}_{R^*} = n^{-1} \sum_{i=1}^{n} \{Y_i - V_i d_i (D^T D)^{-1} D^T (\varphi_1 - \varphi_2)/2\}.$$

Suppose Conditions 1 and 2 hold. Then, as $n \to \infty$ *with* K *fixed, both* $n^{1/2}\hat{\tau}_{R1} - n^{1/2}\hat{\tau}_{R^*}$ *and* $n^{1/2}\hat{\tau}_{R2} - n^{1/2}\hat{\tau}_{R^*}$ *converge in probability to zero. Furthermore,* $n^{1/2}(\hat{\tau}^*_R - \bar{\tau})$ *converges in distribution to a Gaussian random variable with mean zero and variance*

$$\sigma^2_{R^*} = \lim_{n\to\infty} (4n)^{-1} (\varphi_1 - \varphi_2)^T (I - H_D)(\varphi_1 - \varphi_2).$$

Theorem 1 characterizes several useful properties of the regression assisted estimator $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ under Conditions 1 and 2. First, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are consistent estimators of $\bar{\tau}$. Second, $n^{1/2}(\hat{\tau}_{R1} - \bar{\tau})$ and $n^{1/2}(\hat{\tau}_{R2} - \bar{\tau})$ are asymptotically equivalent and have Gaussian limiting distributions. Finally, the asymptotic variance of the regression assisted estimators are at least as small as that of the classical difference in means estimator for a paired experiment. To demonstrate the final point,

$$\begin{aligned}
\text{var}(n^{1/2}\hat{\tau}_{R^*} \mid \mathcal{F}, \mathcal{Z}) &= n^{-1}\text{var}\left(\sum_{i=1}^{n} Z_{i1}\{\varphi_{i1} - \varphi_{i2} - d_i (D^T D)^{-1} D^T (\varphi_1 - \varphi_2)\} \mid \mathcal{F}, \mathcal{Z} \right) \\
&= (4n)^{-1} (\varphi_1 - \varphi_2)^T (I - H_D)(\varphi_1 - \varphi_2) \\
&\leq (4n)^{-1} (\varphi_1 - \varphi_2)^T (\varphi_1 - \varphi_2) = \text{var}(n^{1/2}\hat{\tau}_C \mid \mathcal{F}, \mathcal{Z}),
\end{aligned}$$

where the inequality stems from H_D being a projection matrix (and hence, being positive semi-definite). Importantly, these properties hold regardless of the truth of the regression model yielding the projection matrices H_D and H_A.

Perhaps surprisingly, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ have the same asymptotic variance despite the fact that $\hat{\tau}_{R2}$ is based on a regression which included both the difference between covariates VD and the relative level of the covariates M. For insight into why this is the case, the average value of the covariates in a given matched pair does not vary across randomizations, while the difference between the covariates for the treated and control individuals does, taking values d_i and $-d_i$ with equal probability for each pair i. If the slope coefficients on the level of covariates m_i have limits as $n \to \infty$ (as guaranteed by Conditions 1 and 2), then the contribution of m_i to the prediction of Y_i also does not vary across randomizations, and hence does not contribute to the variance of the estimator. In light of this, is there motivation for including the level of the covariates when estimating $\bar{\tau}$? As we now demonstrate, inclusion of M in the regression allows for the construction of variance estimators that are less conservative then those derived from a regression excluding M, thus allowing for more powerful inference while still maintaining the desired size α.

5 Enabling Neyman-style inference

As described in §3 inference using the classical difference-in-means estimator, $\hat{\tau}_C$, typically proceeds using an upper bound on the variance. If one chooses a consistent estimator of an upper bound, inference performed is then asymptotically conservative. As will be demonstrated, the nominal variance estimators for the intercept coefficients derived from linear model theory with fixed design and homoskedastic Gaussian errors can be employed towards this end.

The standard variance estimators for the intercept coefficients $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ returned by statistical software are

$$S^2_{R1} = (n-K-1)^{-1}\frac{(Y-\hat{\tau}_{R1}e)^T(I-VH_DV)(Y-\hat{\tau}_{R1}e)}{e^T(I-VH_DV)e},$$
$$S^2_{R2} = (n-K-L-1)^{-1}\frac{(Y-\hat{\tau}_{R2}e)^T(I-H_A)(Y-\hat{\tau}_{R2}e)}{e^T(I-H_A)e}.$$

Theorem 2. *Under Conditions 1 and 2, $nS^2_{R1} - var(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F}, \mathcal{Z})$ converges in probability to*

$$\lim_{n\to\infty}(4n)^{-1}(\varphi_1+\varphi_2-2\bar{\tau}e)^T(\varphi_1+\varphi_2-2\bar{\tau}e) \geq 0.$$

Under the same conditions, $nS^2_{R2} - var(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F}, \mathcal{Z})$ converges in probability to

$$\lim_{n\to\infty}(4n)^{-1}(\varphi_1+\varphi_2-2\bar{\tau}e)^T(I-H_M)(\varphi_1+\varphi_2-2\bar{\tau}e) \geq 0.$$

Corollary 1. *Under Conditions 1 and 2, $nS^2_{R1} - nS^2_{R2}$ converges in probability to*

$$\lim_{n\to\infty}(4n)^{-1}(\varphi_1+\varphi_2-2\bar{\tau}e)^T H_M(\varphi_1+\varphi_2-2\bar{\tau}e) \geq 0.$$

Under an additive treatment effect model both S^2_{R1} and S^2_{R2} are consistent for $\text{var}(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F}, \mathcal{Z})$ itself. Otherwise, S^2_{R1} is a consistent upper bound under Conditions 1 and 2. S^2_{R2} can also be consistent for $\text{var}(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F}, \mathcal{Z})$ if the relative level of the covariates in a pair m_i is perfectly predictive of the aggregate treatment effect in pair i, $\tau_{i1}+\tau_{i2}$; otherwise, it too is a consistent upper bound. As a consequence, $1-\alpha$ confidence intervals for $\bar{\tau}$ of the form $\hat{\tau}_{R1} \pm \Phi^{-1}(1-\alpha/2)S_{R1}$ and $\hat{\tau}_{R2} \pm \Phi^{-1}(1-\alpha/2)S_{R2}$ will be asymptotically conservative as desired. Intervals constructed using S_{R2} will be shorter than those constructed using S_{R1} unless the level of the covariates M is not predictive of the average value of the treatment effects in pairs. Asymptotically, inference for $\bar{\tau}$ using $\hat{\tau}_{R2}$ and S^2_{R2} will never be less powerful than that performed using $\hat{\tau}_{R1}$ and S^2_{R1} or $\hat{\tau}_C$ and S^2_C and will maintain the desired size α.

6 A simulation study

6.1 The general setup

In this simulation study, we assess the finite sample properties of both estimation and inference using $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$. In any given simulation scenario, we generate S experiments, each consisting of n pairs. In the sth of S experiments, the values of the $2n$ pairs potential outcomes

$\{(r_{Tij}, r_{Cij}), i = 1, .., n, j = 1, 2\}$ are generated dependent on (potentially unknown) covariates w_{ij} and noise terms ε_{Tij} and ε_{Cij} as

$$r_{Tij} = \mu_T(w_{ij}) + \varepsilon_{Tij}; \quad r_{Cij} = \mu_C(w_{ij}) + \varepsilon_{Cij}$$

for some functions $\mu_T(\cdot)$ and $\mu_C(\cdot)$. The average treatment effect for experiment s is

$$\bar{\tau}^{(s)} = (2n)^{-1} \sum_{i=1}^{n} \sum_{j=1}^{2} \{\mu_T(w_{ij}) - \mu_C(w_{ij}) + \varepsilon_{Tij} - \varepsilon_{Cij}\},$$

and the goal in each experiment s is to estimate $\bar{\tau}^{(s)}$ and create confidence intervals for $\bar{\tau}^{(s)}$. We additionally simulate observed covariates x_{ij} which need not equal w_{ij}. The researcher then stipulates functions $f(\cdot)$ and $g(\cdot)$ of the observed covariates, creates $d_i = f(x_{i1}) - f(x_{i2})$ and $m_i = \{g(x_{i1}) + g(x_{i2})\}/2 - n^{-1} \sum_{\ell=1}^{n} \{g(x_{\ell 1}) + g(x_{\ell 2})\}/2$, and creates D with d_i^T as the ith row and M with m_i^T as the ith row. These values are fixed across randomizations within experiment s.

In the bth of B randomizations for the sth experiment, we assign the treatment to exactly one unit in each pair. Afterwards, we calculate $\hat{\tau}_C$, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ based on the observed responses. Further, we construct the corresponding 95% confidence intervals for the experiment-specific average treatment effect $\bar{\tau}^{(s)}$ using S_C^2, S_{R1}^2, and S_{R2}^2. Through our B randomizations, we estimate the true coverage of nominally 95% confidence intervals, the relative magnitude of the standard errors which are used to construct confidence intervals, and relative root mean square errors for the three estimators over the randomization distribution for the sth experiment.

6.2 Linear regression under a nonlinear truth

Our study contains two simulation settings inspired by the setting presented in Kang and Schafer (2007). We simulate four covariates w_{ijp}, $p = 1, .., 4$ that are unknown to the researcher. For each p, we stipulate that w_{i1p} is Gaussian distributed with mean zero and variance 1, w_{i2p} Gaussian distributed with mean w_{i1p} and variance 1/4, and w_{ijp} is independent of $w_{ijp'}$ for $p' \neq p$. We further set $\varepsilon_{Tij} = \varepsilon_{Cij}$, simulate ε_{Tij} from a Gaussian distribution with mean 0 and variance 1, and set $f(x_{ij}) = g(x_{ij}) = x_{ij}$. The revealed covariates x_{ij} for each individual are complicated functions of w_{ij},

$$x_{ij1} = \exp(w_{ij1}/2); \quad x_{ij2} = w_{ij2}/\{1 + \exp(w_{ij1})\} + 10$$
$$x_{ij3} = \{(w_{ij1} w_{ij3})/25 + 0.6\}^3; \quad x_{ij4} = (w_{ij2} + w_{ij4} + 20)^2.$$

The settings vary in the functions $\mu_T(\cdot)$ and $\mu_C(\cdot)$. The two possibilities for $\mu_T(\cdot)$ and $\mu_C(\cdot)$ are (1) parallel response surfaces with $\mu_T(w_{ij}) = \mu_C(w_{ij}) = 27.4 w_{ij1} + 13.7(w_{ij2} + w_{ij3} + w_{ij4})$; and (2) nonparallel response surfaces with $\mu_T(w_{ij}) = 27.4 w_{ij1} + 13.7(w_{ij2} + w_{ij3} + w_{ij4})$ and $\mu_C(w_{ij}) = 13.7(w_{ij1} + w_{ij2}) + 3 w_{ij3} + 27.4 w_{ij4}$

As the researcher runs a multiple regression using the observed covariates x_{ij} instead of w_{ij}, the researcher is fitting an linear model using covariates that have a highly nonlinear relationship with the potential outcomes. Recall that the results in Theorems 1 and 2 are robust to the truth of the regression model, and instead leverage the adjustment as an algorithmic fit. We will now assess how well the asymptotics apply in small and moderately sized paired experiments.

Table 1: Classical and regression assisted average treatment effect estimators with a misspecified response function

	True Coverage of 95% CIs			100 × Stand. Error Ratios			100 × RMSE Ratios	
	C	$R1$	$R2$	$R1:C$	$R2:C$	$R2:R1$	$R1:C$	$R2:C$
Parallel								
$n=25$	95 (94, 96)	95 (94, 96)	94 (89, 97)	55 (37, 77)	55 (37, 77)	100 (91, 107)	54 (37, 78)	55 (38, 85)
$n=50$	95 (94, 96)	95 (94, 96)	94 (92, 96)	54 (41, 69)	53 (41, 67)	99 (94, 101)	54 (41, 69)	55 (41, 72)
$n=100$	95 (94, 96)	95 (94, 96)	95 (93, 96)	54 (45, 65)	54 (44, 64)	99 (96, 100)	54 (44, 65)	55 (45, 67)
$n=500$	95 (94, 96)	95 (94, 96)	95 (94, 96)	54 (50, 60)	54 (50, 60)	100 (99, 100)	54 (50, 61)	55 (50, 61)
Nonparallel								
$n=25$	100 (99, 100)	100 (99, 100)	99 (96, 100)	90 (77, 90)	54 (40, 71)	60 (43, 83)	81 (55, 125)	63 (43, 98)
$n=50$	100 (99, 100)	100 (99, 100)	100 (99, 100)	86 (78, 94)	53 (43, 63)	61 (50, 74)	64 (48, 88)	53 (39, 74)
$n=100$	100 (100, 100)	100 (100, 100)	100 (100, 100)	86 (80, 90)	53 (46, 60)	62 (53, 71)	55 (45, 69)	48 (39, 61)
$n=500$	100 (100, 100)	100 (100, 100)	100 (100, 100)	85 (82, 97)	53 (50, 56)	62 (53, 71)	45 (40, 51)	44 (39, 50)

6.3 Results

In each setting we simulate $S = 1000$ experiments which contain n matched pairs. Within each experiment, we conduct $B = 1000$ randomizations. We conducted our simulation study with experiments of size $n = 25, 50, 100$ and $n = 500$. Given the experiment sizes considered, the extent to which the results presented in this work are applicable for experiments of these sizes is far but certain at the onset of this investigation.

Table 1 shows the result of our simulation study. The first three columns characterize the actual coverage properties of 95% confidence intervals for the sample average treatment effect constructed using the standard formulae for the standard errors of $\hat{\tau}_C$, $\hat{\tau}_{R1}$, and $\hat{\tau}_{R2}$ respectively. The second three columns compare the relative values of these standard errors based on these three estimators (and hence, the relative width of confidence intervals across experiments). The final two column shows the relative values for the true root mean squared errors of the regression estimators $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ relative to the classical estimator $\hat{\tau}_C$. The initial number provided in each entry is the median value across the 1000 experiments, and the endpoints of the interval provided in parenthesis are the 2.5% and 97.5% quantiles across the 1000 experiments.

In the parallel response setting with $\mu_T(\cdot) = \mu_C(\cdot)$, the performed inference will be asymptotically exact rather conservative since we set $\varepsilon_{Tij} = \varepsilon_{Cij}$ implying an additive treatment effect

model. As Table 1 indicates, for all values of n the median coverage of intervals constructed using $\hat{\tau}_C$ and $\hat{\tau}_{R1}$ for $\bar{\tau}^{(s)}$ over the $S = 1000$ experiments was 95%, with coverage falling in the range (94%, 96%) for 95% of the simulated experiments. For $\hat{\tau}_{R2}$, the median coverage was 94% at $n = 25$, with a 95% range of (89%, 97%) indicating a more drastic deviation from claimed coverage across randomizations in some experiments for small n. The reason for this is not small sample bias in the estimator $\hat{\tau}_{R2}$ itself, but rather a systematic downward bias of S_{R2} relative to the true value of $\text{var}(\hat{\tau}_{R2} \mid \mathcal{F}, \mathcal{Z})^{1/2}$. In fact, the 95% range for $E(S_{R2}|\mathcal{F}, \mathcal{Z})/\text{var}(\hat{\tau}_{R2} \mid \mathcal{F}, \mathcal{Z})^{1/2}$ across experiments was (0.81, 1.05) with $n = 25$. The coverage improves quickly as n increases, and presents little reason for alarm once $n = 100$. The median ratio of both the standard errors and the root mean squared errors from $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ relative to $\hat{\tau}_C$ are slightly over $1/2$ across experiments, indicating the potential benefits from estimation and inference using regression adjustment in this setting. As predicted by Theorem 2, the ratios of the standard error estimates based on $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are clustered about 1 due to the treatment effect being additive, indicating that even in finite samples there is nothing to be gained from using $\hat{\tau}_{R2}$ over $\hat{\tau}_{R1}$ under additivity.

For the settings with a nonparallel response surface, intervals constructed will be conservative since the treatment effect is no longer additive. Indeed, we see that for all values of n the median coverage across experiments was either 99% or 100% for confidence intervals based on all three estimators whose nominal coverage was 95%. In these settings, we also observe the potential benefits of using the estimator $\hat{\tau}_{R2}$. For $\hat{\tau}_{R1}$, the standard errors used to construct confidence intervals were only slightly smaller than those constructed using $\hat{\tau}_C$, with the median ratio of the two across experiments decreasing from 0.9 for $n = 25$ to 0.85 at $n = 500$. On the other hand, the median ratio of interval widths constructed using $\hat{\tau}_{R2}$ relative to those using $\hat{\tau}_C$ was slightly over $1/2$, while the median ratio relative to those constructed based on $\hat{\tau}_{R1}$ was 0.6. To emphasize, this result does not stem from an inherently lower variance for the estimator $\hat{\tau}_{R2}$, as reflected in the ratios of the true root mean square errors for estimation using $\hat{\tau}_C$ relative to $\hat{\tau}_{R1}$ and relative to $\hat{\tau}_{R2}$ shown in Table 1. Indeed, the mean squared errors for $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are equal in the limit by Theorem 1. Rather, as Corollary 1 describes, this stems from S^2_{R2} being a tighter upper bound for $\text{var}(\hat{\tau}_{R2} \mid \mathcal{F}, \mathcal{Z})$ than S^2_{R1} is for $\text{var}(\hat{\tau}_{R1} \mid \mathcal{F}, \mathcal{Z})$. Hence, intervals based on $\hat{\tau}_{R2}$ are shorter than those constructed using $\hat{\tau}_{R1}$ while maintaining the proper coverage.

For $n = 25$, confidence intervals constructed using $\hat{\tau}_{R2}$ had coverage that fell short of their nominal coverage in certain experiments in the parallel response setting. In this case, variants of sandwich estimators that target improved small-sample performance may provide intervals with better coverage. Indeed, rerunning the simulation at $n = 25$ using the standard error estimator of Efron (1982) (commonly referred to as 'HC3' in the literature and in statistical software) which is recommended by Long and Ervin (2000) for use in small samples, the median coverage of nominally 95% confidence intervals for $n = 25$ with parallel response surfaces was 98% with a 95% range of (96%, 100%), indicating a slight over correction but nonetheless yielding conservative inference. In the nonparallel setting intervals remained conservative at $n = 25$, and the gains from using S_{R2} relative to S_{R1} persisted but to a lesser extent, with a median ratio of 0.86 and a 95% interval of (0.60, 1.25). Use of 'HC3' thus conferred some of the benefits of using $\hat{\tau}_{R2}$ as an estimator in the presence of nonparallel responses while maintaining the desired size of the performed inference. For the simulations with $n = 50$, 100, and 500, the results of Theorem 1, Theorem 2, and Corollary 1 seem to yield an accurate representation of finite sample

performance of estimation and inference using $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ along with the conventional standard errors S_{R1} and S_{R2}.

A Facts and Lemmas

The following fact about the covariates M will be useful moving forward.

Fact 1. *$e^T m_\ell = 0$ and $(1 - H_M)e = e$ due to centering of the covariates m_ℓ, $\ell = 1, ..., L$*

We now prove three lemmas which will be utilized in the proofs of our main results.

Lemma 1. *$n^{-1}A^TA$ converges in probability to a block-diagonal matrix with $\lim_{\to\infty} n^{-1}D^TD$ in the upper-left block of size $K \times K$, $\lim_{n\to\infty} n^{-1}M^TM$ in the lower-right block of size $L \times L$, and zeroes in the remaining entries.*

Proof. It suffices to show that $n^{-1}\sum_{i=1}^n V_i d_{ik} m_{i\ell}$ converges in probability to zero for any $k = 1, ..., K$, $\ell = 1, ..., L$. Trivially, $E(n^{-1}\sum_{i=1}^n V_i d_{ik} m_{i\ell} \mid \mathcal{F}, \mathcal{Z}) = 0$. Through Chebyshev's inequality, we can show that $\text{var}(n^{-1}\sum_{i=1}^n V_i d_{ik} m_{i\ell} \mid \mathcal{F}, \mathcal{Z}) \to 0$ to complete the proof. The proof utilizes Condition 1, along with the Cauchy-Schwarz inequality.

$$\begin{aligned}
\text{var}\left(n^{-1}\sum_{i=1}^n V_i d_{ik} m_{i\ell} \mid \mathcal{F}, \mathcal{Z}\right) &= n^{-2}\sum_{i=1}^n (d_{ik})^2 (m_{i\ell}^2) \\
&\leq n^{-2}\left(\sum_{i=1}^n d_{ik}^4\right)^{1/2}\left(\sum_{i=1}^n m_{ik}^4\right)^{1/2} \\
&\leq C/n
\end{aligned}$$

by Condition 1, which tends to zero as $n \to \infty$.

Lemma 2. *$n^{-1}\sum_{i=1}^n V_i Y_i d_{ik}$ converges in probability to $\lim_{n\to\infty}(2n)^{-1}\sum_{i=1}^n(\varphi_{i1} - \varphi_{i2})d_{ik}$ for any $k = 1, ..., K$. Similarly, $n^{-1}\sum_{i=1}^n Y_i m_{ik}$ converges in probability to $\lim_{n\to\infty}(2n)^{-1}\sum_{i=1}^n(\varphi_{i1} + \varphi_{i2})m_{ik}$ for any $\ell = 1, ..., L$*

Proof. We prove the result for $n^{-1}\sum_{i=1}^n V_i Y_i d_{ik}$; the proof for $n^{-1}\sum_{i=1}^n Y_i d_{ik}$ is analogous. First, $E(n^{-1}\sum_{i=1}^n V_i Y_i d_{ik} \mid \mathcal{F}, \mathcal{Z}) = (2n)^{-1}(\varphi_{i1} - \varphi_{i2})d_{ik}$. We now show that $\text{var}(n^{-1}\sum_{i=1}^n V_i Y_i d_{ik} \mid \mathcal{F}, \mathcal{Z})$ converges to zero.

$$\begin{aligned}
\text{var}\left(n^{-1}\sum_{i=1}^n V_i Y_i d_{ik} \mid \mathcal{F}, \mathcal{Z}\right) &= (2n)^{-2}\sum_{i=1}^n(\varphi_{i1} + \varphi_{i2})^2 (d_{ik})^2 \\
&\leq (2n)^{-2}\left\{\sum_{i=1}^n(\varphi_{i1} - \varphi_{i2})^4\right\}^{1/2}\left(\sum_{i=1}^n d_{ik}^4\right)^{1/2} \\
&\leq C/n
\end{aligned}$$

by Condition 1, which tends to zero as $n \to \infty$. Through Chebyshev's inequality, we then have that $n^{-1}\sum_{i=1}^{n} V_i Y_i d_{ik}$ converges in probability to $\lim_{n\to\infty}(2n)^{-1}\sum_{i=1}^{n}(\varphi_{i1}-\varphi_{i2})d_{ik}$, where the limit in the right-hand side exists through Condition 2.

Lemma 3. *$n^{-1}e^T(I - VH_DV)e$ and $n^{-1}e^T(I - H_A)e$ converge in probability to zero.*

Proof. We prove the result for $n^{-1}e^T(I - H_A)e$, and in the process we also prove the result for $n^{-1}e^T(I - VH_DV)e$.

$$\begin{aligned} n^{-1}e^T(I - H_A)e &= 1 - n^{-1}\sum_{i=1}^{n} a_i(A^TA)^{-1}A^Te \\ &= 1 - n^{-1}\sum_{i=1}^{n} V_i d_i \left\{(D^TD)^{-1}D^TVe\right\} + o_p(1) \\ &= n^{-1}e^T(I - VH_DV)e + o_p(1), \end{aligned}$$

where the $o_p(1)$ term comes from Fact 1, Condition 2, and Lemma 1.

Note that $E(V_iV_j \mid \mathcal{F}, \mathcal{Z}) = 0$ for $i \neq j$ due to independence of matched pairs i and j. Let h_{ij} be the i, j entry of H_D, and recall that $\sum_{i=1}^{n} h_{ii} = K$ and $0 \leq h_{ii} \leq 1$ since H_D is a projection matrix, and that $h_{ii}(1 - h_{ii}) = \sum_{j\neq i} h_{ij}^2$ due to idempotence of H_D. We then have $E(n^{-1}e^T(I - VH_DV)e^T \mid \mathcal{F}, \mathcal{Z}) = (n - K)/n \to 1$. Through Markov's inequality, it suffices to show that $\mathrm{var}(n^{-1}e^T(I - VH_DV)e \mid \mathcal{F}, \mathcal{Z})$ converges to zero.

$$\begin{aligned} \mathrm{var}(n^{-1}e^T(I - VH_DV)e \mid \mathcal{F}, \mathcal{Z}) &= n^{-2}\sum_{j\neq i} 2h_{ij}^2 \\ &= n^{-2}\sum_{i=1}^{n} 2h_{ii}(1 - h_{ii}) \leq n^{-2}(2K), \end{aligned}$$

which converges to zero as desired.

B Proof of Theorem 1

We prove the result for $n^{1/2}\hat{\tau}_{R2}$. Define $\hat{\beta}_D$ and $\hat{\beta}_M$ as

$$(\hat{\beta}_D, \hat{\beta}_M) = (A^TA)^{-1}A^TY.$$

That is, $\hat{\beta}_D$ and $\hat{\beta}_M$ are vectors of length K and L corresponding to the sample slopes for the covariates VD and M respectively from a regression of Y on A. Define β_D and β_M as

$$\begin{aligned} \beta_D &= (D^TD)^{-1}D^T(\varphi_1 - \varphi_2)/2 \\ \beta_M &= (M^TM)^{-1}M^T(\varphi_1 + \varphi_2)/2, \end{aligned}$$

and note further that $\sum_{i=1}^{n} m_{ik}\hat{\beta}_k = 0$ and $\sum_{i=1}^{n} m_{ik}\beta_k = 0$ for $k = 1, .., K$ by Fact 1.

$n^{1/2}\hat{\tau}_{R2}$ can be written in the following form:

$$n^{1/2}\hat{\tau}_{R2} = \frac{n^{-1/2}\sum_{i=1}^{n}(Y_i - V_i d_i \hat{\beta}_D)}{n^{-1}e^T(I - H_A)e}$$

By Fact 1, Lemma 1 and Condition 2

$$\sum_{i=1}^{n} n^{-1/2} V_i d_i \hat{\beta}_D = n^{-1/2}\sum_{i=1}^{n} V_i d_i \{(n^{-1}D^T D)^{-1} n^{-1} D^T V Y + o_p(1)\}.$$

By Lemma 3, the denominator of $n^{1/2}\hat{\tau}_{R2}$ is $1 + o_p(1)$. Hence,

$$\begin{aligned} n^{1/2}\hat{\tau}_{R2} - n^{1/2}\hat{\tau}_{R^*} = &-n^{1/2}\frac{\sum_{i=1}^{n} V_i d_i}{n}\left\{(n^{-1}D^T D)^{-1} n^{-1} D^T V Y - \beta_D + o_p(1)\right\} \\ &+ n^{1/2}\frac{\sum_{i=1}^{n} Y_i}{n}\{o_p(1)\} \end{aligned}$$

By Lemma 2 and Condition 2, $(n^{-1}D^T D)^{-1} n^{-1} D^T V Y - \beta_D$ is $o_p(1)$. By Condition 2, the finite-sample central limit theorem of Hájek et al. (1999, §V.1.2, p. 153) applies to $n^{1/2}\sum_{i=1}^{n} V_i d_{ik}/n$ for $k = 1, .., K$ and to $n^{1/2}\sum_{i=1}^{n} Y_i/n$, and hence these terms are $O_p(1)$. Therefore, $n^{1/2}\hat{\tau}_{R2} - n^{1/2}\hat{\tau}_{R^*}$ converge in probability to zero as desired. Asymptotic normality of $n^{1/2}\hat{\tau}_{R^*}$ follows through the special central limit theorem of Hájek et al. (1999, §V.1.2, p. 153) by Condition 2.

To prove the result for $\hat{\tau}_{R1}$, simply redefine $\hat{\beta}_D$ by $(D^T D)^{-1} D^T V Y$ and disregard the mention of β_M and $\hat{\beta}_M$.

C Proof of Theorem 2

Recall the definitions for S^2_{R1} and S^2_{R2}:

$$\begin{aligned} S^2_{R1} &:= (n - K - 1)^{-1}\frac{(Y - \hat{\tau}_{R1}e)^T(I - VH_D V)(Y - \hat{\tau}_{R1}e)}{e^T(I - VH_D V)e}, \\ S^2_{R2} &:= (n - K - L - 1)^{-1}\frac{(Y - \hat{\tau}_{R2}e)^T(I - H_A)(Y - \hat{\tau}_{R2}e)}{e^T(I - H_A)e}. \end{aligned}$$

We prove the remark for S^2_{R2}. Note first that by Lemma 3, $n\{e^T(I - H_A)e\}^{-1}$ converges in probability to 1. Re-expressing S^2_{R2} using standard identities for residuals from linear regression

$$\begin{aligned} nS^2_{R2} = &\, n^{-1}\{1 + o_p(1)\}(Y - \hat{\tau}_{R2}e)^T(I - H_M)(Y - \hat{\tau}_{R2}e) &(1)\\ &- n^{-1}\{1 + o_p(1)\}(Y - \hat{\tau}_{R2}e)^T(I - H_M)VD\{D^T V(I - H_M)VD\}^{-1}DV(I - H_M)(Y - \hat{\tau}_{R2}e) &(2) \end{aligned}$$

We begin by assessing term (1). By Condition 2, Lemma 2 and Theorem 1, (1) converges in probability to

$$\lim_{n\to\infty}(2n)^{-1}\sum_{i=1}^{n}\{(\varphi_{i1} - \bar{\tau})^2 + (\varphi_{i2} - \bar{\tau})^2)\} - \lim_{n\to\infty}(4n)^{-1}(\varphi_1 + \varphi_2 - 2\bar{\tau}e)^T H_M(\varphi_1 + \varphi_2 - 2\bar{\tau}e)$$

We now assess term (2). Recall that any term of the form $n^{-1}\sum_{i=1}^{n} V_i d_{ik} m_{i\ell}$ or $n^{-1}M^T V D$ converges in probability to zero as described in the proof of Lemma 1, and recall that $H_M = M(M^TM)^{-1}M^T$. Hence, by Condition 2, Lemma 2, and Theorem 1, (2) converges in probability to

$$-\lim_{n\to\infty}(4n)^{-1}(\varphi_1-\varphi_2)^T H_D(\varphi_1-\varphi_2).$$

Meanwhile, $\text{var}(n^{1/2}\hat{\tau}_{R2} \mid \mathcal{F}, \mathcal{Z}) \to \lim_{n\to\infty}(4n)^{-1}(\varphi_1-\varphi_2)(I-H_D)(\varphi_1-\varphi_2)$. Subtracting these two quantities, $nS_{R2}^2 - \text{var}(n^{1/2}\hat{\tau}_{R2} \mid \mathcal{F}, \mathcal{Z})$ converges in probability to

$$\begin{aligned}
&\lim_{n\to\infty}\left[(2n)^{-1}\sum_{i=1}^{n}\{(\varphi_{i1}-\bar{\tau})^2+(\varphi_{i2}-\bar{\tau})^2\} - (4n)^{-1}(\varphi_1+\varphi_2-2\bar{\tau}e)^T H_M(\varphi_1+\varphi_2-2\bar{\tau}e)\right.\\
&\left.-(4n)^{-1}\sum_{i=1}^{n}(\varphi_{i1}-\varphi_{i2})^2\right]\\
&= \lim_{n\to\infty}(4n)^{-1}\left(\sum_{i=1}^{n}(\varphi_{i1}+\varphi_{i2}-2\bar{\tau})^2 - (\varphi_1+\varphi_2-2\bar{\tau}e)^T H_M(\varphi_1+\varphi_2-2\bar{\tau}e)\right)\\
&= \lim_{n\to\infty}(4n)^{-1}(\varphi_1+\varphi_2-2\bar{\tau}e)^T(I-H_M)(\varphi_1+\varphi_2-2\bar{\tau}e)
\end{aligned}$$

The proof for S_{R1}^2 follows by replacing $I-H_M$ with I and $\hat{\tau}_{R2}$ with $\hat{\tau}_{R1}$ in (1) and (2) above.

References

Aronow, P. M. and Middleton, J. A. (2013). A class of unbiased estimators of the average treatment effect in randomized experiments. *Journal of Causal Inference*, 1(1):135–154.

Bloniarz, A., Liu, H., Zhang, C.-H., Sekhon, J., and Yu, B. (2016). Lasso adjustments of treatment effect estimates in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(27):7383–7390.

Ding, P. (2016). A paradox from randomization-based causal inference. *Statistical Science*, to appear.

Efron, B. (1982). *The jackknife, the bootstrap and other resampling plans*. SIAM, Philadelphia, PA.

Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40(2):180–193.

Greevy, R., Lu, B., Silber, J. H., and Rosenbaum, P. (2004). Optimal multivariate matching before randomization. *Biostatistics*, 5(2):263–275.

Hájek, J., Sĭdák, Z., and Sen, P. K. (1999). *Theory of Rank Tests*. Academic press.

Hodges, J. L. and Lehmann, E. L. (1963). Estimates of location based on rank tests. *The Annals of Mathematical Statistics*, 34(2):598–611.

Imai, K. (2008). Variance identification and efficiency analysis in randomized experiments under the matched-pair design. *Statistics in Medicine*, 27(24):4857–4873.

Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Kang, J. D. and Schafer, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science*, pages 523–539.

Lin, W. (2013). Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics*, 7(1):295–318.

Long, J. S. and Ervin, L. H. (2000). Using heteroscedasticity consistent standard errors in the linear regression model. *The American Statistician*, 54(3):217–224.

Neyman, J. (1923). On the application of probability theory to agricultural experiments. Essay on principles. Section 9 (in Polish). *Roczniki Nauk Roiniczych*, X:1–51. Reprinted in Statistical Science, 1990, 5(4):463-480.

Rosenbaum, P. R. (1993). Hodges-Lehmann point estimates of treatment effect in observational studies. *Journal of the American Statistical Association*, 88(424):1250–1253.

Rosenbaum, P. R. (2002). Covariance adjustment in randomized experiments and observational studies. *Statistical Science*, 17(3):286–327.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701.